BakBone

Whitepaper

NetVault 6.5.x Virtual Disk Library Backup Staging Guide

Rev 2
October 1, 2002

Table of Contents

Click on the content and it will automatically link you to that section of the white paper.

Introduction

This document describes the NetVault Virtual Disk Library (VDL) feature. It provides information on what it is, how it works, when to use it, how to set up a VDL, and an example of how to configure a Full, Incremental and Consolidated backup job that uses VDL as a target.

What is a Virtual Disk Library (VDL)

A VDL gives disk storage the look and feel of a tape library including advanced media management features with the added benefit of speed and reduced cost associated with disk. Gain more manageability by setting specific backup policies such as retention dates, rotation schemes, media groups. The media manager does not distinguish between VDL "Tape" Libraries red resources. Save sets can be accessed wherever they reside. Send Incrementals to disk for ultra fast restores from disk. Tape now becomes a strategic component of a data protection strategy, rather than the primary ingredient. Use NetVault to create multiple duplications of backup jobs from the Virtual Disk Library, to tape, or vice versa. By storing backup data on a VDL, you can run data copy or duplication jobs off line, without impacting your network, application servers or workstations. With Virtual Disk Libraries, there are none of the mechanical penalties of tape backup over a network, such as shoe-shining, or a slow data stream host. Virtual Disk Libraries captures the data, whether its trickled, blasted, or gently passed along, it all ends up at the receiving end of Virtual Media Slots as a saveset, improving backup window performance markedly.

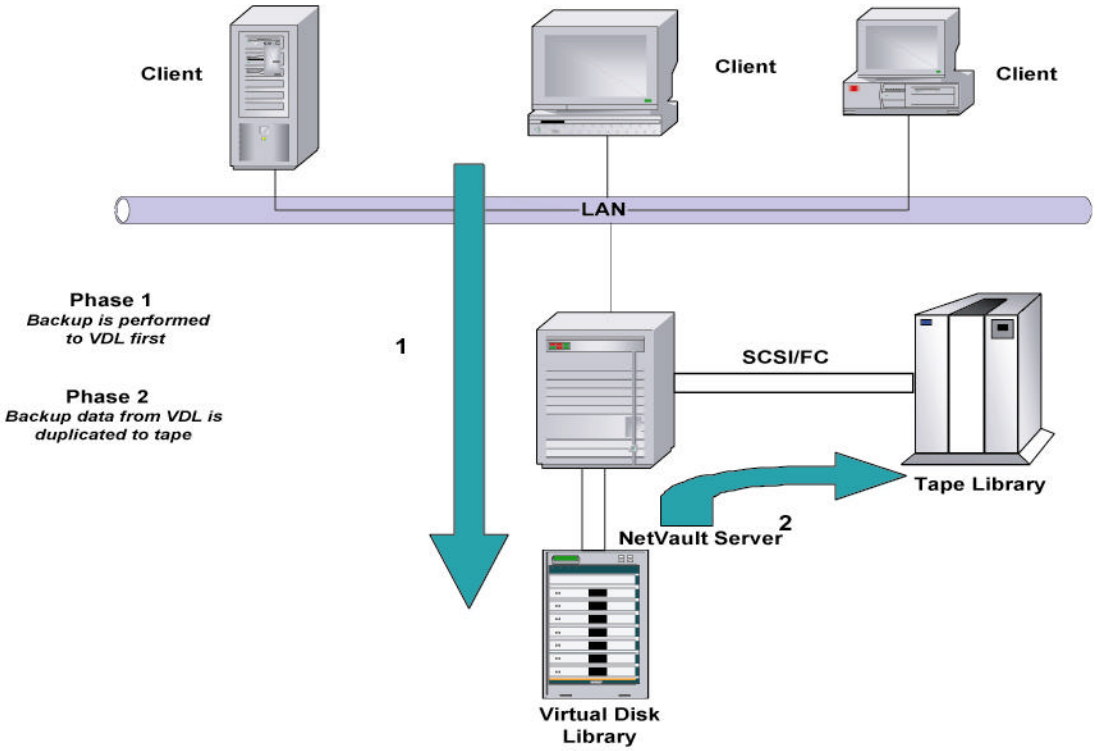

During the first phase of a staged job, all selected data is backed up to the VDL (disk). Once this process completes, the second phase know as a *Data Copy* begins and copies the saveset from the virtual library to physical tape. This two-phase process is automatic and is defined in the job. After the data copy completes, the space used in the VDL is cleared and available for another backup.

When to use a VDL

There are two reasons why you might want to use VDL Staging:

- When backing up a file system that contains many (millions) of files, the server may simply not be able to read the files quickly enough to stream today's high-performance tape drives. This will result in shoe-shining, which may lead to premature drive or media failures. It is always critical to attempt to effectively stream tape hardware at its native rates. This shoe-shine condition does not apply to disk storage, so there is no penalty from slow performance when backing up to a VDL.

- If the backup window is too small to backup several Clients onto a limited number of tape drives, then create a VDL with enough Virtual Drives to backup all Clients simultaneously. The performance depends on the network bandwidth, a gigabit network would almost surely be required.

  For example: 5 Clients need to backup in one hour. Each Client has 10GB of data and there is one DLT7000 tape drive which is capable of 18GB/hr. Obviously, the backup window won't be met because there is a 50GB of data and the drive is only capable of 18GB/hour. The solution is to use Disk Staging and backup to multiple Virtual Tape Drives first and then copy to the real physical tape. Define enough Virtual tape drives to complete backups of all Clients within one hour. Later, the copy operation may take up to 3 hours to complete. Note: This happens off-line on the backup server only and will not affect the Clients= or network.

VDL Staging vs. Multiplexing

NetVault deals with the problem of multiple Clients backing up to limited tape drives in short backup windows by using, staging. Other backup products may choose to resolve this problem by doing what is commonly called multiplexing. With Multiplexing, you can have multiple streams of backup data can be sent to one tape device. There are many disadvantages to the Mulitplexing approach:

- A backup of any given client will *span* much more tape than is actually required. This results in having to handle multiple tapes per client backup.

- Leads to a higher probability of *failure* if one of the media goes bad since more media is used for any given backup.

- Leads to *much longer* restore times because much more tape needs to be scanned for a given restore time. This is due to the fact that the data must be reconstructed from the multiple data streams and this takes time.

- Uses much more CPU time on the backup server because the data streams must be reorganized and packed into a multiplexed stream. This could cause *severe performance* problems with today's high-speed tape devices.
- Requires a *proprietary* tape format so the competitor's backup products would have to be reinstalled in a disaster recovery situation to restore the data.

With the NetVault Staging approach, backups of each client are always contiguous on tape thus using less tape and making restore times quick. A non-proprietary tape format is also used. The only requirement is that extra disk space is required for the virtual library resource allocation.

The NetVault Virtual Disk Library (VDL)

A virtual disk library is a directory structure on disk in which there are directories called *drives* and *slots*. Beneath each of these directories are numbered directories. Each numbered directory defines a unique slot or drive.

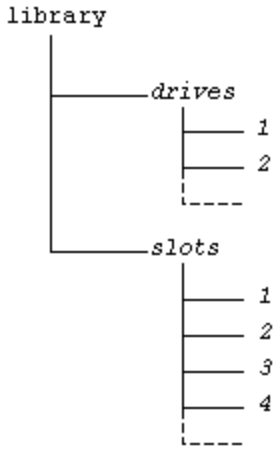

Within each slot numbered directory resides a *media* file of a specified size. These are the virtual library's "tapes". NetVault treats virtual libraries just like real physical libraries in every way and these are even licensed as libraries.

The more drives the VDL contains, the more simultaneous backups can be performed.

There *must* always be as many or more *slots* as there are *drives* for the library to function properly. It is recommended that the virtual library be configured with a minimum of 8 slots. By having extra slots, it allows NetVault to properly handle backup retention cycles. Also, different operating systems impose limits on the maximum size a file (media) may be which will affect the number of slots required.

At configuration time, when defining the quantity of slots and media capacity, NetVault will create media files of the given size in MB and pre-allocate the space.

Virtual Disk Library Strategies

Some of the most effective uses of a VDL are in conjunction with NetVault's Consolidated File System backup feature. The Consolidated File System Backup, starts with one full backup and uses the incremental/differential backups, which follow to create a new full. The results of a consolidated backup is equivalent to a full backup performed at the same time as the last incremental.

Typical Consolidated File System Backup schedule:

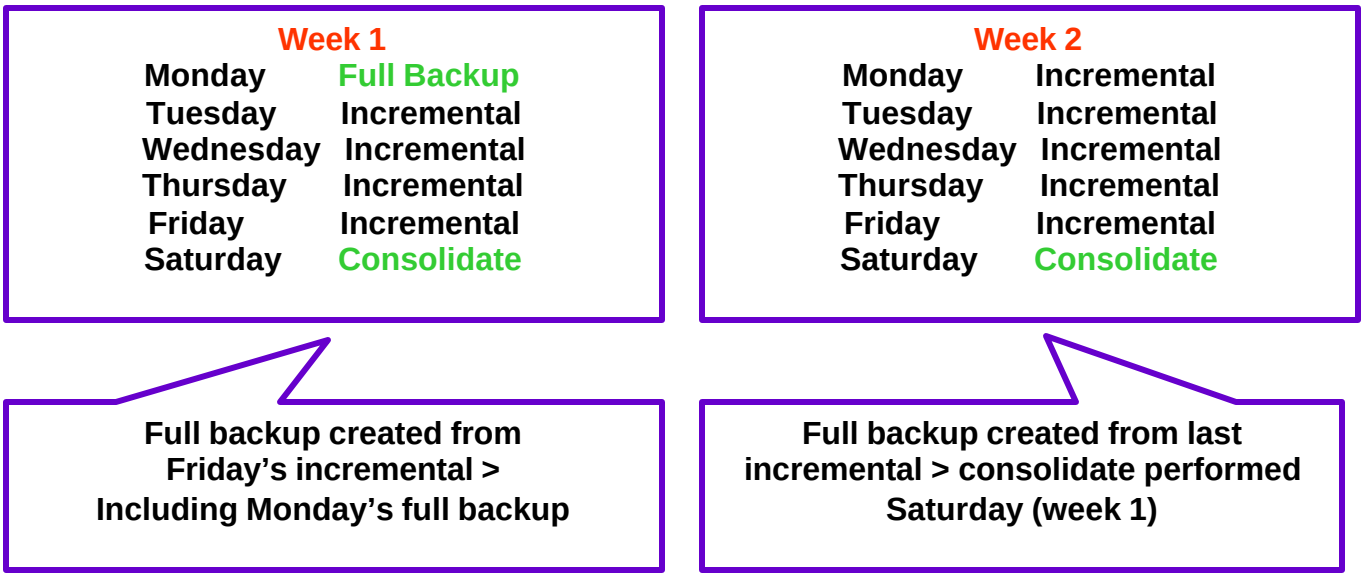

NetVault can mix VDL's and tape as backup media, which allows for a variety of different strategies for:

- Increased protection of data
- Added backup speed
- Added restore speed
- Reduced data vulnerability through multiply copies

Examples:

Consolidating Incrementals:

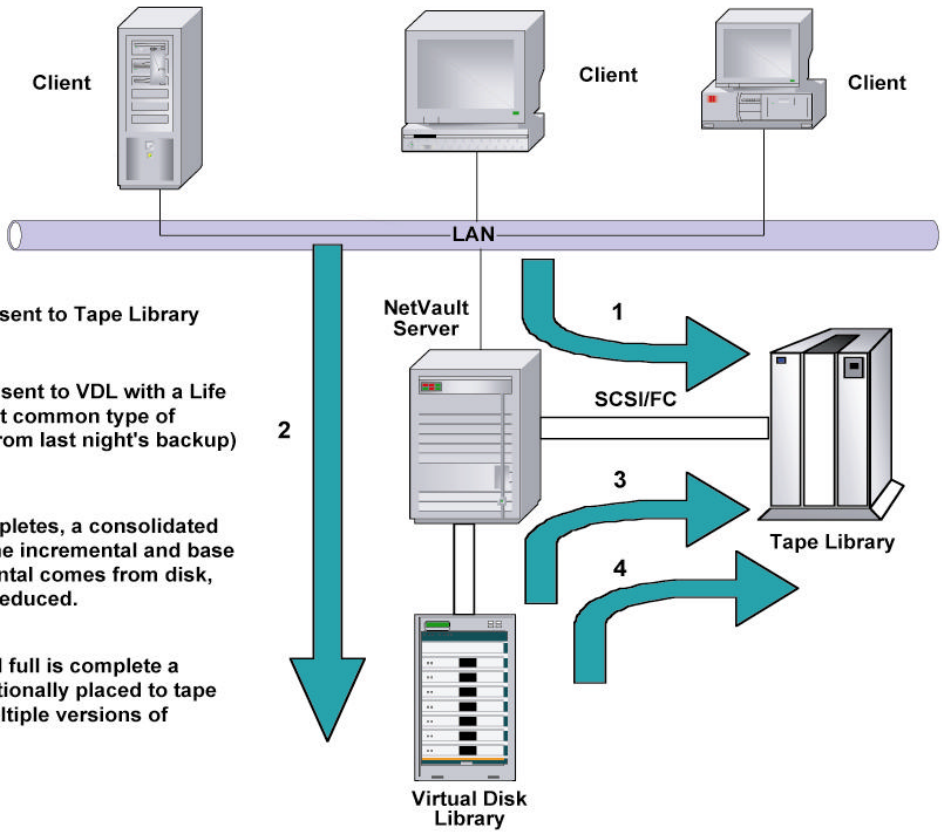

Step 1: Base full backups are sent to Tape Library as Normal

Step 2: Nightly Incremental is sent to VDL with a Life of 3 days. This allows for most common type of restore (recover a single file from last night's backup) to come from VDL.

Step 3: After incremental completes, a consolidated full is created/updated from the incremental and base line full. Because the incremental comes from disk, consolidation time is greatly reduced.

Step 4: After the consolidated full is complete a copy of the incremental is optionally placed to tape so the customer can have multiple versions of files.

Application recovery Strategy:

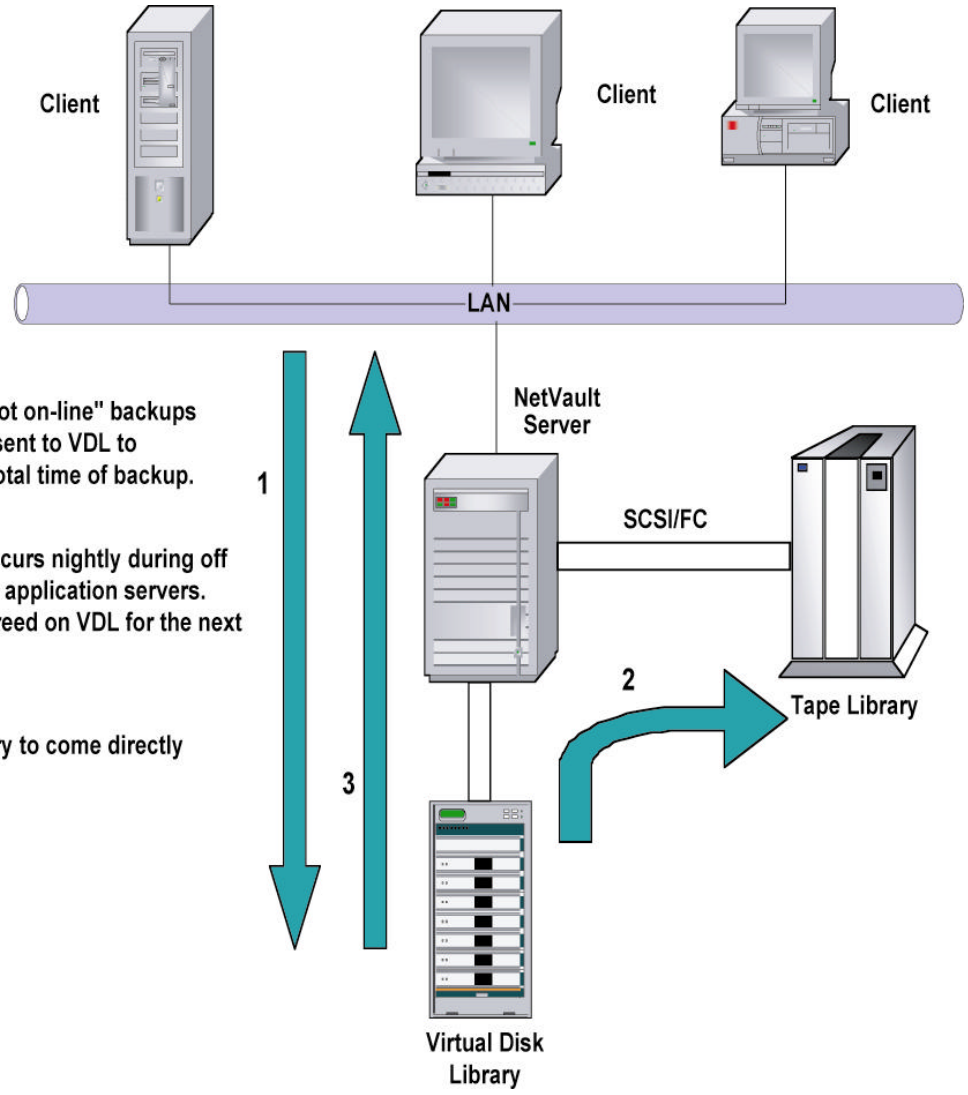

Step 1: Critical application "hot on-line" backups are done during the day and sent to VDL to maximize speed and lessen total time of backup.

Step 2: Duplication to tape occurs nightly during off hours, without interruption to application servers. Space is then automatically freed on VDL for the next days backups.

Step 3:This allows for recovery to come directly from the VDL.

Getting Started

Determine the VDL size, configuration and location:

- Decide how big the library needs to be. If you are meeting a tight backup window and have several clients and you must backup simultaneously, the VDL must be big enough to accommodate the total amount of data of all the clients. If on the other hand, you have the problem of a server with a very large file system, you may only require enough space in your VDL to handle one backup at a time.

- Decide on your VDL geometry with regards to drives and slots. The number of simultaneous backup jobs you need to run will dictate the number of drives. The number of slots will be dictated by the total size you require and your operating system. It is recommended that you have a minimum of 8 slots. Make sure you know your available free disk space before you attempt to create the virtual library, as NetVault will pre-allocate the space it requires.

- Decide on where to create your VDL. Typically this must be done on the NetVault Server itself or Smart Client Attached Storage.

Also, you must configure your actual physical library or tape drive and have it tested and ready before you can proceed with VDL Staging tests. Refer to the NetVault documentation or QuickStart guide for information on how to do this, as it is not covered in this document.

Installing the Virtual Library Generator

The virtual library generator comes with NetVault installation but is not automatically installed. Install it first before creating the library. Do this as follows:

Step 1

From the main screen, invoke the *Client Management* screen. Right-click on the NetVault Server or the Client installing the Virtual Library generator on and select Install Software.
Navigate the Open file dialog box to the *<netvault_home>/packages/extra-npks* directory and select file called **ddv*nnnn*.npk**.

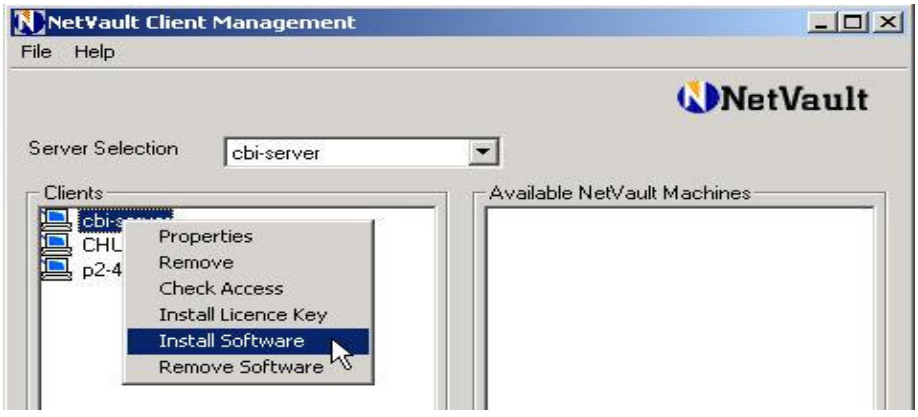

Step 2
The generator is installed.

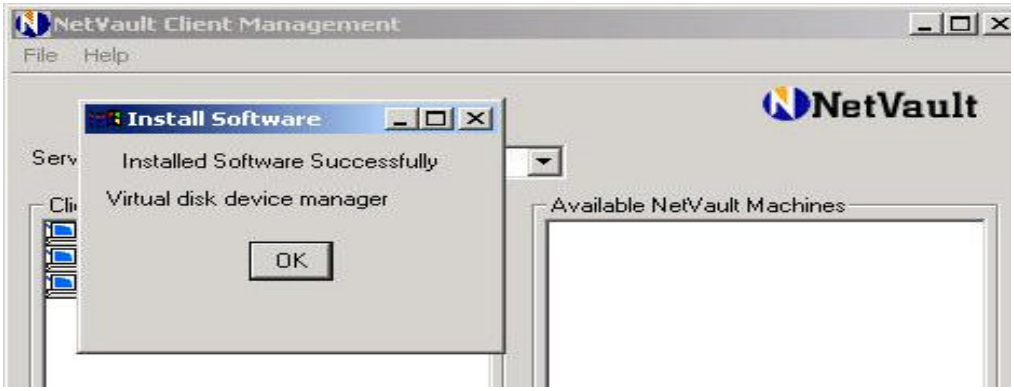

Creating and Configuring the Virtual Library

Follow these steps to create and configure the Virtual Library.

Step 1
Go to the NetVault Device Management screen and select *Add* > *Add Library*.

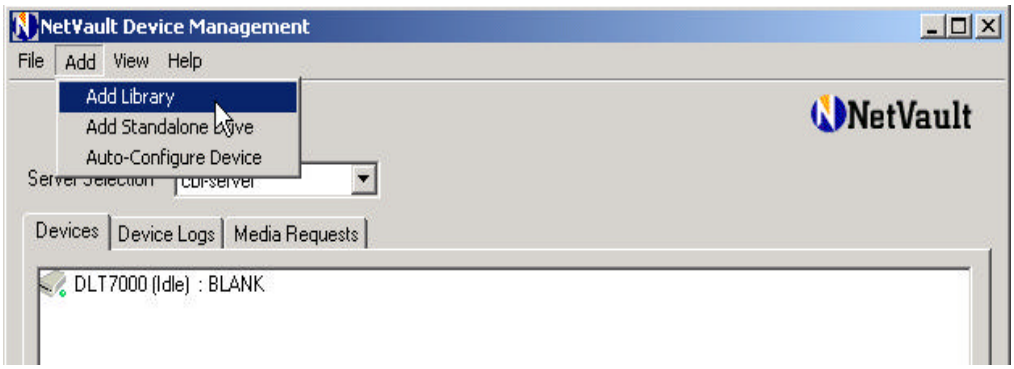

Step 2
Select the server/Smart Client where you want to create the virtual library and right-click. Select *Create virtual library*.

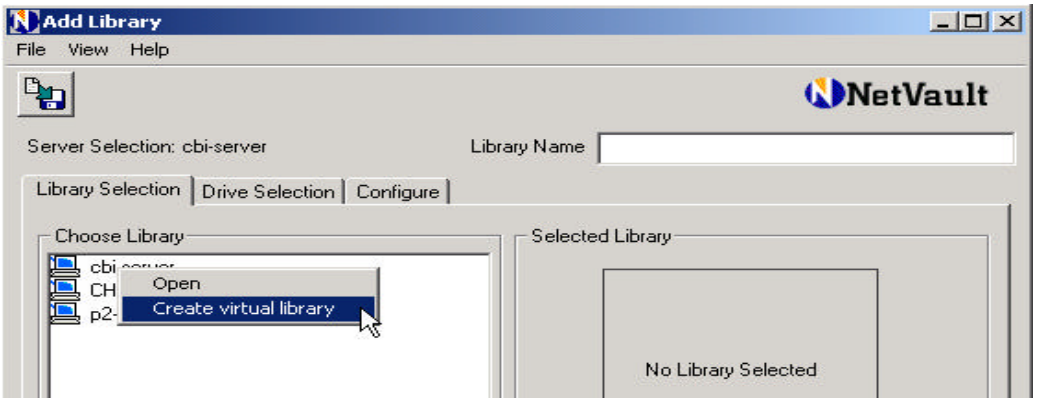

Step 3
The Add new virtual library screen will appear. Enter the *Library Location* on disk where the library will be created. Enter the *Library name* (directory name). Enter the *Number of drives* and *slots* you want. And, enter the *Media capacity* for the media in each slot.

Note: Media capacity should be a minimum of 30MB.

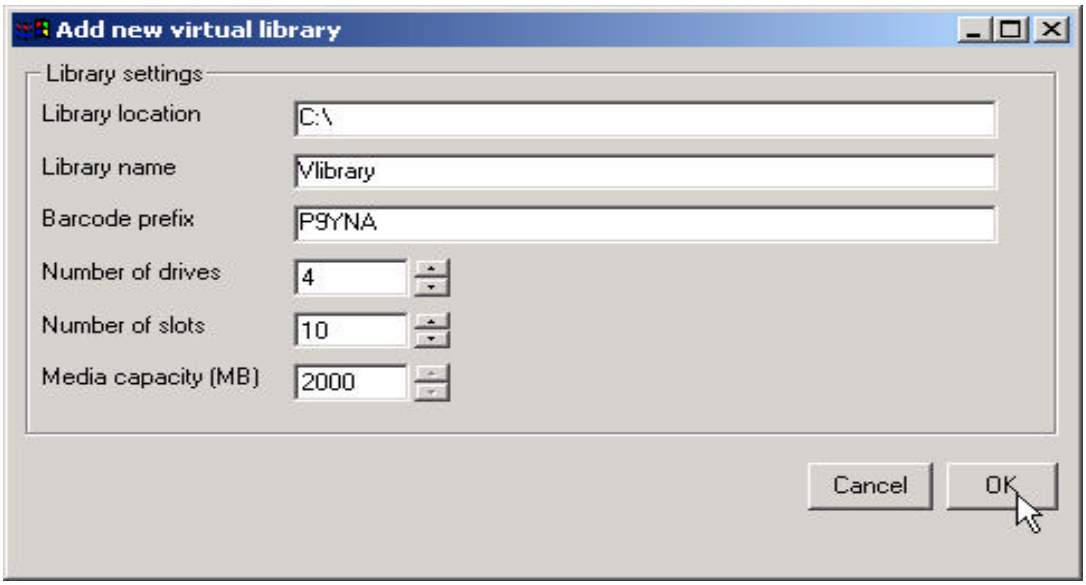

Click OK and the library and empty media files will be created. If creating a large capacity library, this may take a while.

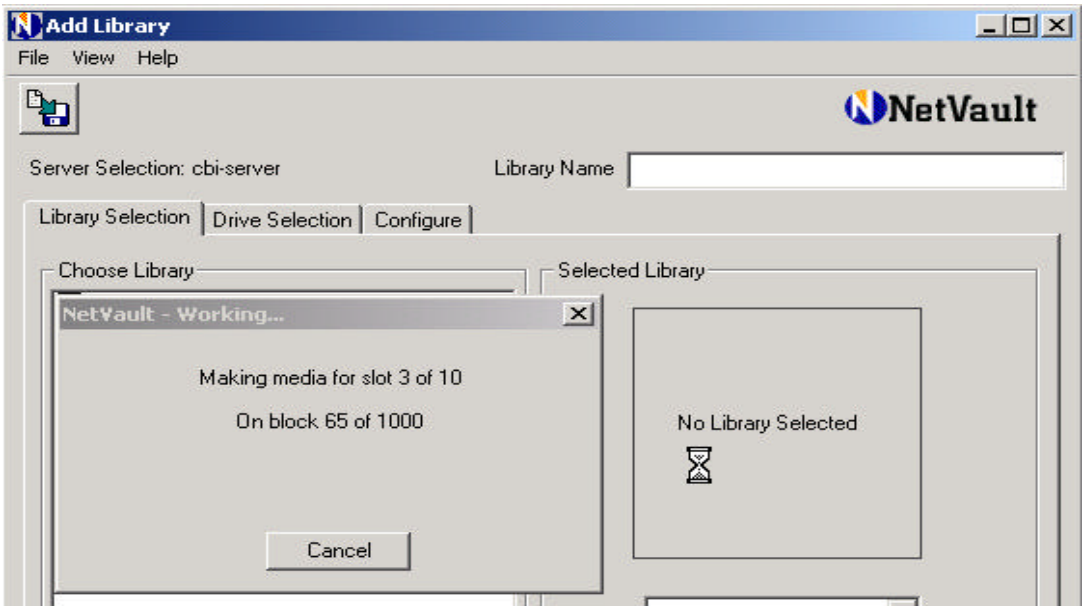

Step 4
Now add the virtual library as you would a real tape library.

In the Library Selection Tab, double-click on the name of the Server on which you created the library. After a scan of your system, a list of all libraries found will be presented. Under Robotic Libraries, double-click on your *virtual library* name.

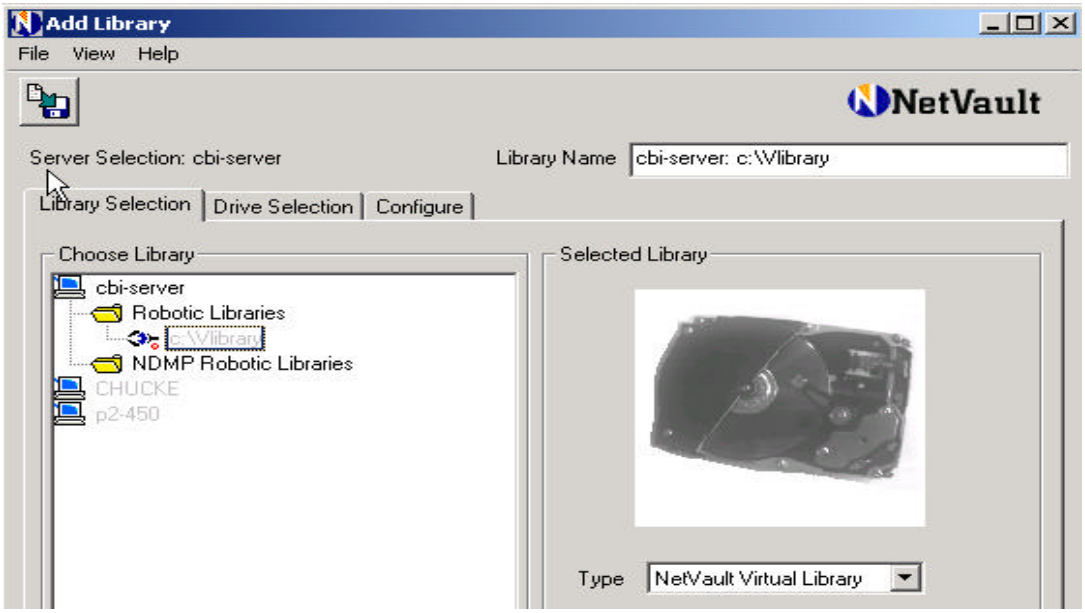

Step 5
Go to the *Drive Selection* Tab. Double-click on the name of the Server on which you created the library. After a scan of your system, a list of all drives will be listed.

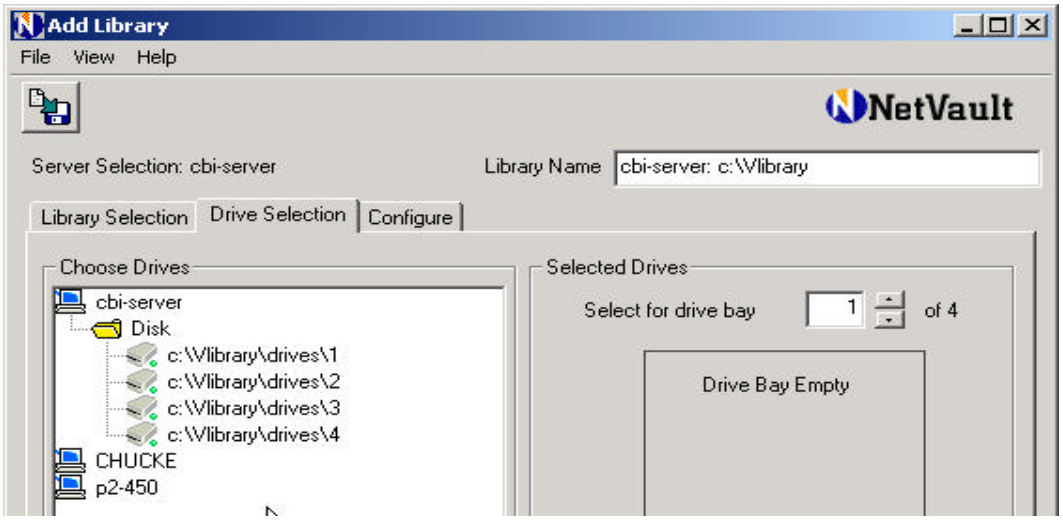

Step 6
With the *Select for drive bay* set to ***1***, double-click on the *first drive* in the list. You now associated virtual drive 1 with bay 1 of the virtual library.

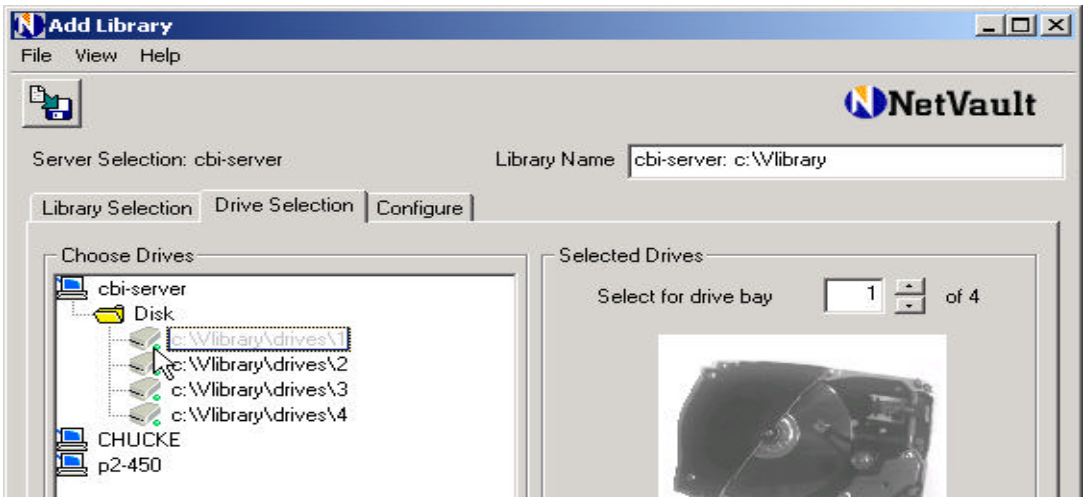

Step 7
With the *Select for drive bay* set to ***2***, double-click on the *second drive* in the list. You now associated virtual drive 2 with bay 2 of the virtual library.

Repeat this step for each drive. Always remember to increase the bay number for each successive drive.

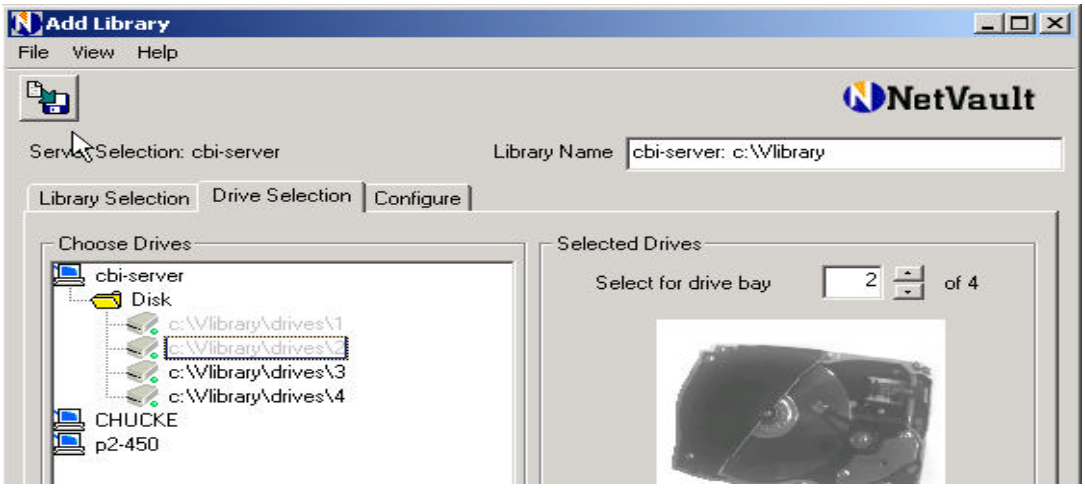

Step 8 (Options)
NetVault allow for several options during the configuration. Two of the most important which apply are.

- Software Compression: Will compress the data written on the VDL drives, but you *must* have the processing power in your system to handle multiple data streams, to take advantage of this feature.
- TurboVault: Allocates main memory for transfer buffers to assist in higher transfer rate for both virtual and physical tape devices when performing local backups.
  (See 6.5.x, Administrator Guide, Appendix A, for further information)

Option 1: Software Compression
Select the *Configure* Tab. Choose *Configure* from the pop-up menu for Drive Bay 1. The Configure Device, box will appear.

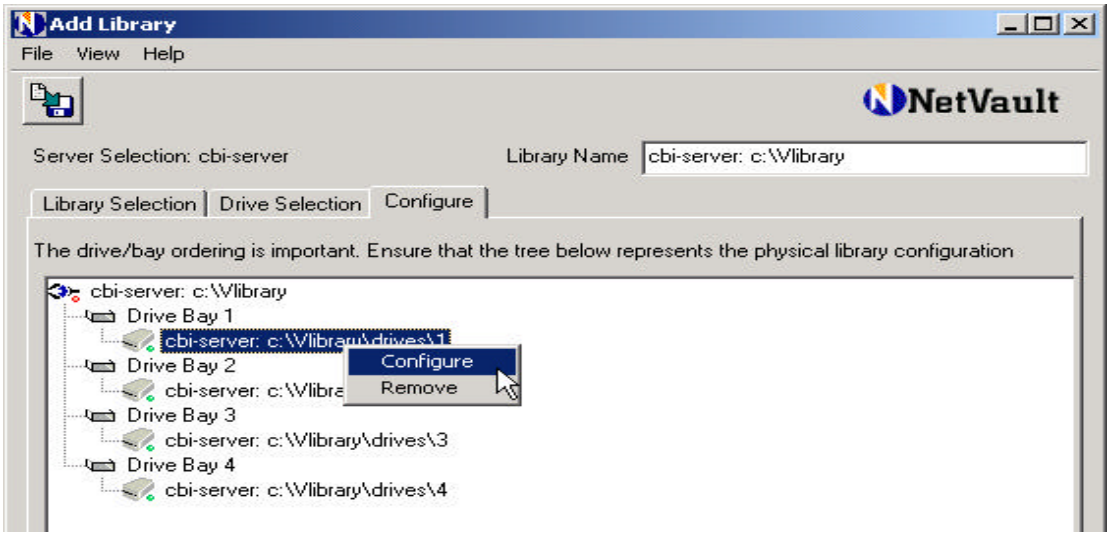

Option 2: TurboVault is located in the *Performance* tab.

Enter an amount (8193) of Main Memory that you want to allocate for data buffering.
8MB of memory is a good starting point, but the correct number for the best performance will depend, on the performance of the disk storage device where the library resides.
(See 6.5.x, Administrator Guide, Appendix A, for further information)
Click ok. Repeat these steps for each *Drive bay* in your VDL

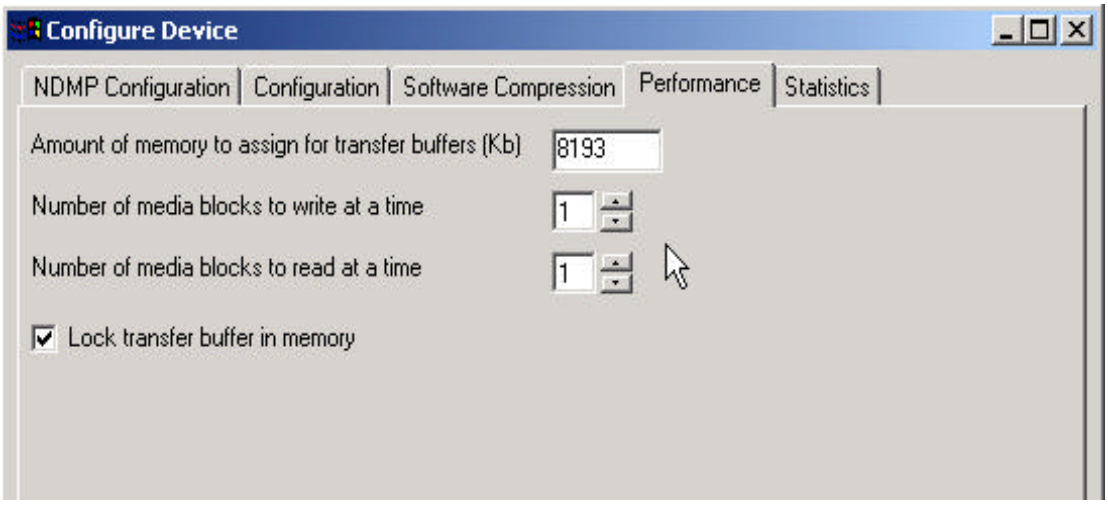

Step 9

Enter a short (one word) name for your library in the *Library Name* field. Click on the *Save details* button on the top left to save the new virtual library.
Close the Add Library screen.

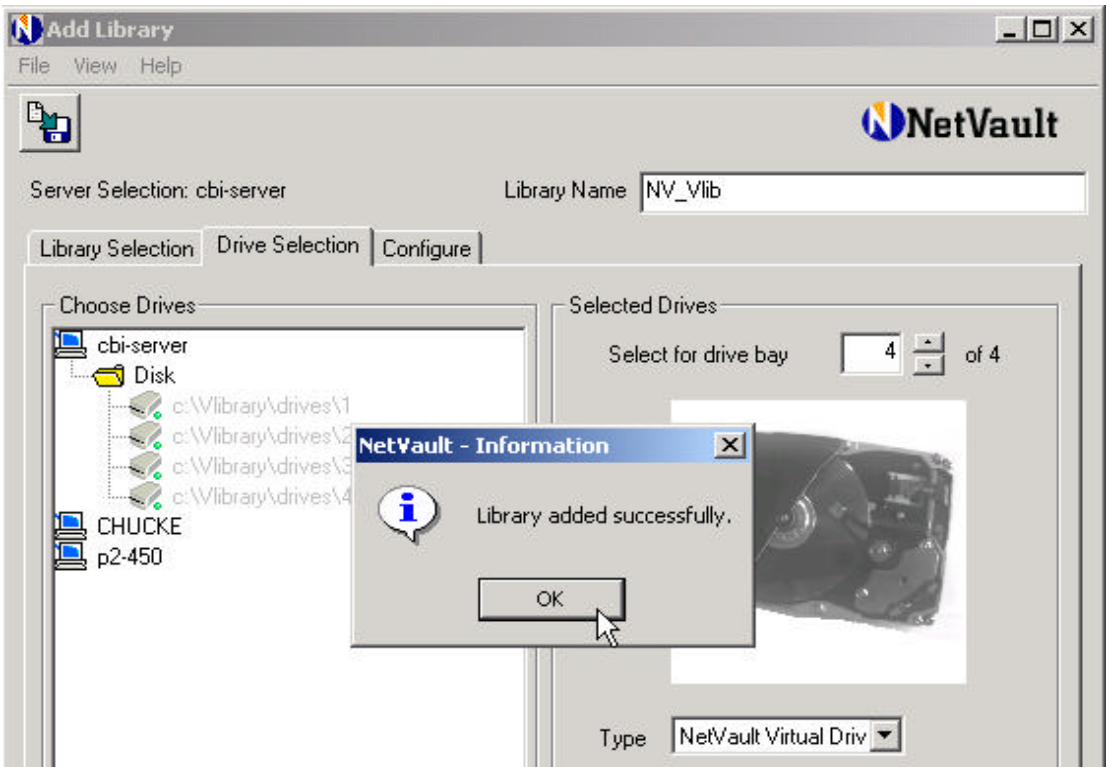

Step 10
The Device Management Window will now display the virtual library. All the media will be displayed as BLANK, which means they are available for backups.

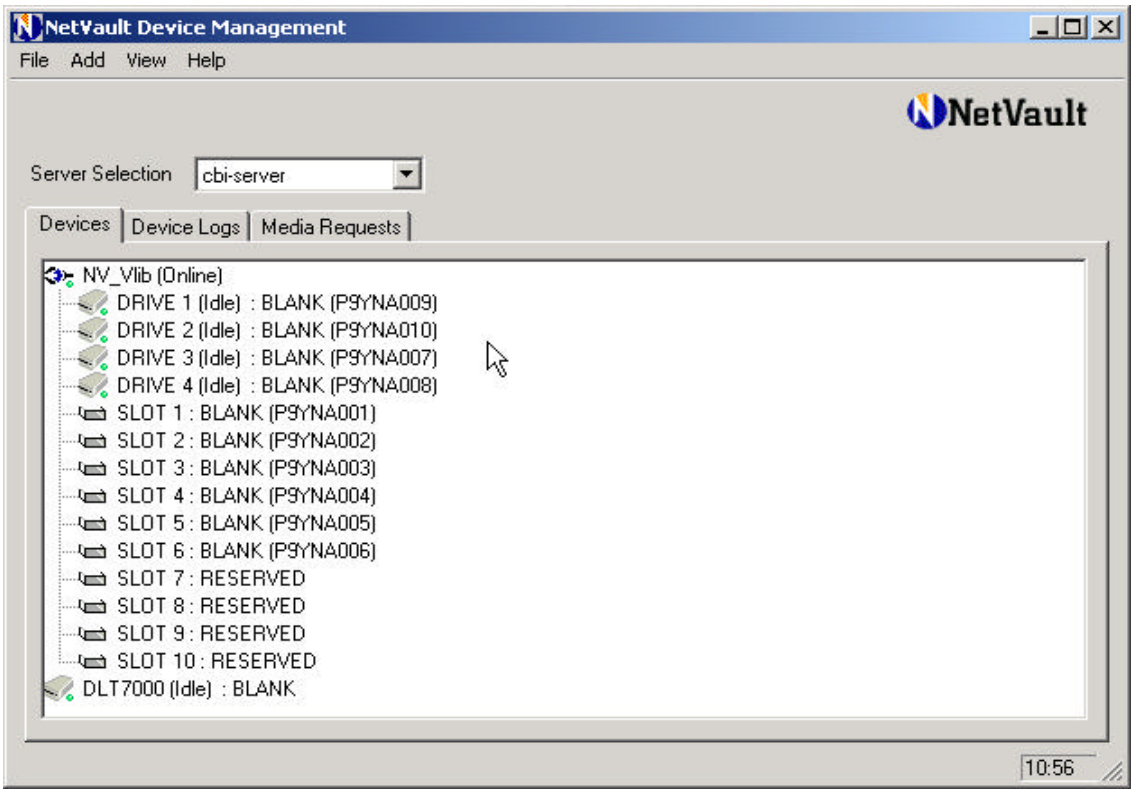

Creating Templates for (VDL) Staging jobs

A VDL Staging job completes in two phases. The first phase backup up the Client data onto the VDL. The second phase copies the backup data from the VDL to a physical tape drive/library. Jobs can be created to automatically perform both phases without user intervention.

The first step is to create three *Target* templates. The first template should target the initial FULL backup operation to the Virtual Library. The second template is for the INCREMENTAL backup target. The third will target the physical drive/library.

NOTE: Besides assigning a target device, Target templates also define other settings such as the very important Grouping and Retention definitions. In production environments you'll probably find that you will have several Target templates defined instead of the three we will use in our example.

Creating the three Target templates

Step 1
From the Main screen, select *Administration > Policy Management.*
The Policy Management screen will appear. *Select the Target Tab*.

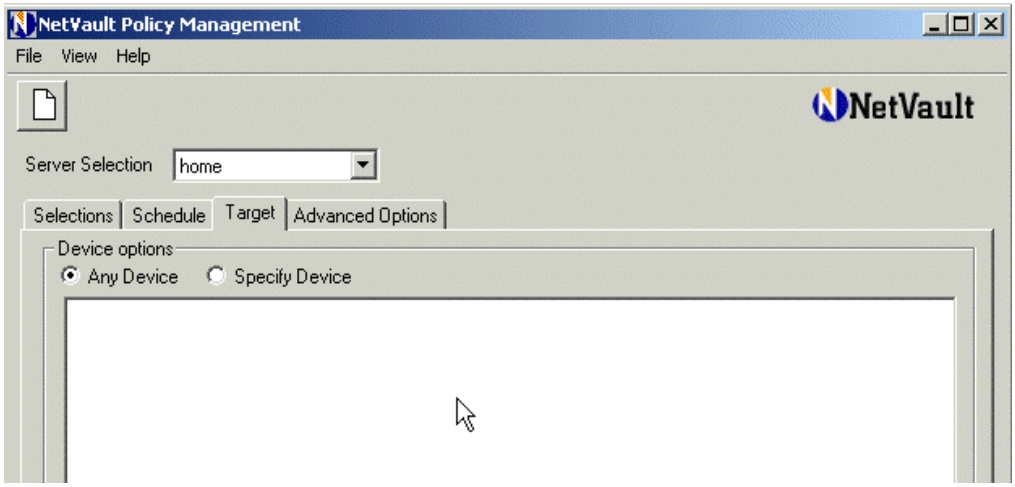

Step 2
Click on *Specify Device* and a list of all drives will be presented. Your screen will look different from the example shown but you will see your Virtual Library drives and your real physical tape resources.

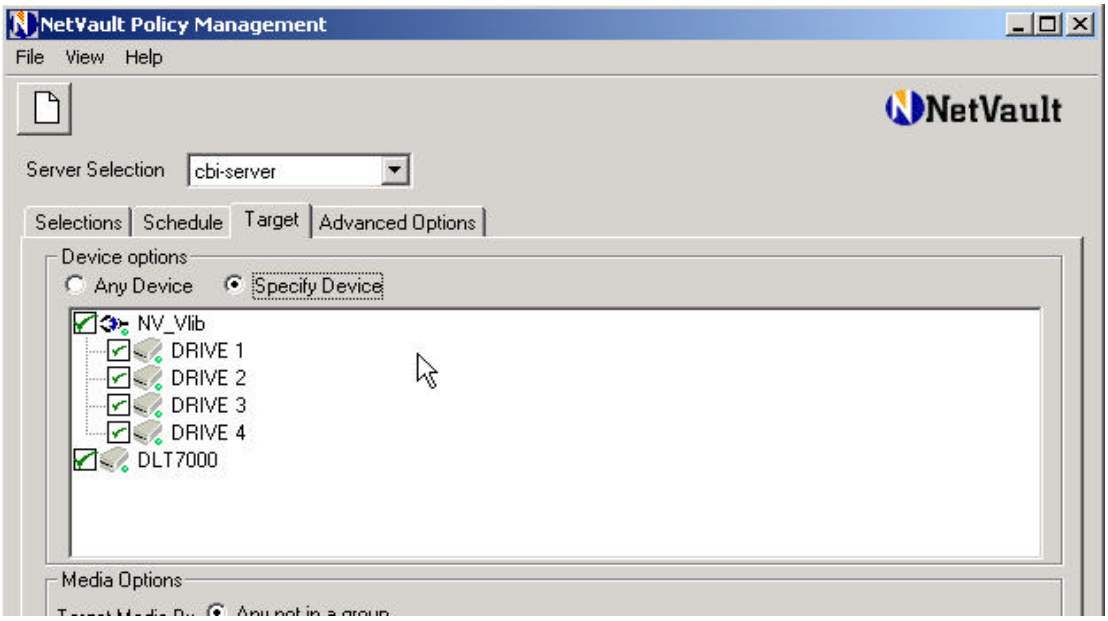

Step 3

The first Target template we will create will send data to the Virtual drives for a FULL backup.

1) De-select all physical drives leaving only the virtual drives ticked.
2) To create separate pools of media, enter a *Group Label*, for example: V_Lib_FULL.
3) As with any job, first select Reuse Media – *Any* in the Media Options.
4) Click on *Save As* at the bottom of the screen. Save the template with a descriptive name such as *Virtual Device FULL Group*.

The template for the FULL to VDL is now ready.

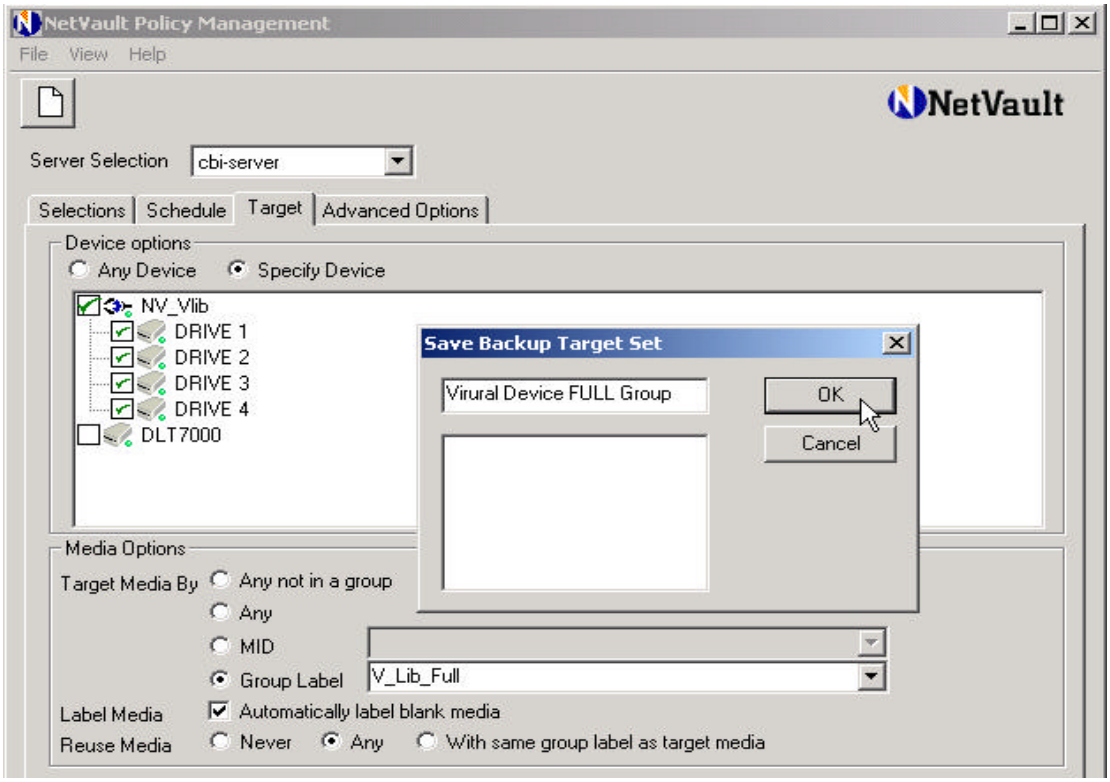

Step 4

Press on the Clear button and again select the Target Tab for our next template.

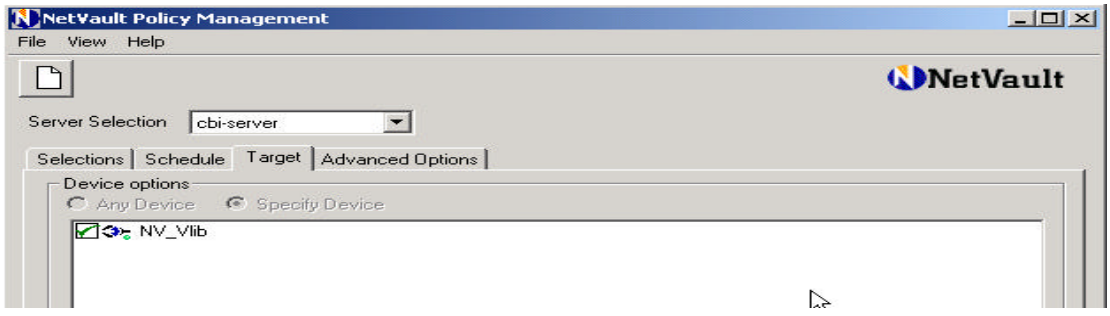

Step 5
The second Target template we will create will send data to the Virtual drives for an Incremental backup.
1) De-select all physical drives leaving only the virtual drives ticked.
2) To create separate pools of media, select *Group Label*, V_Lib_INC.
3) As with any job, select Reuse Media – *Any* in the Media Options.
4) Click on *Save As* at the bottom of the screen. Save the template with a descriptive name such as *Virtual Device INC Group*.

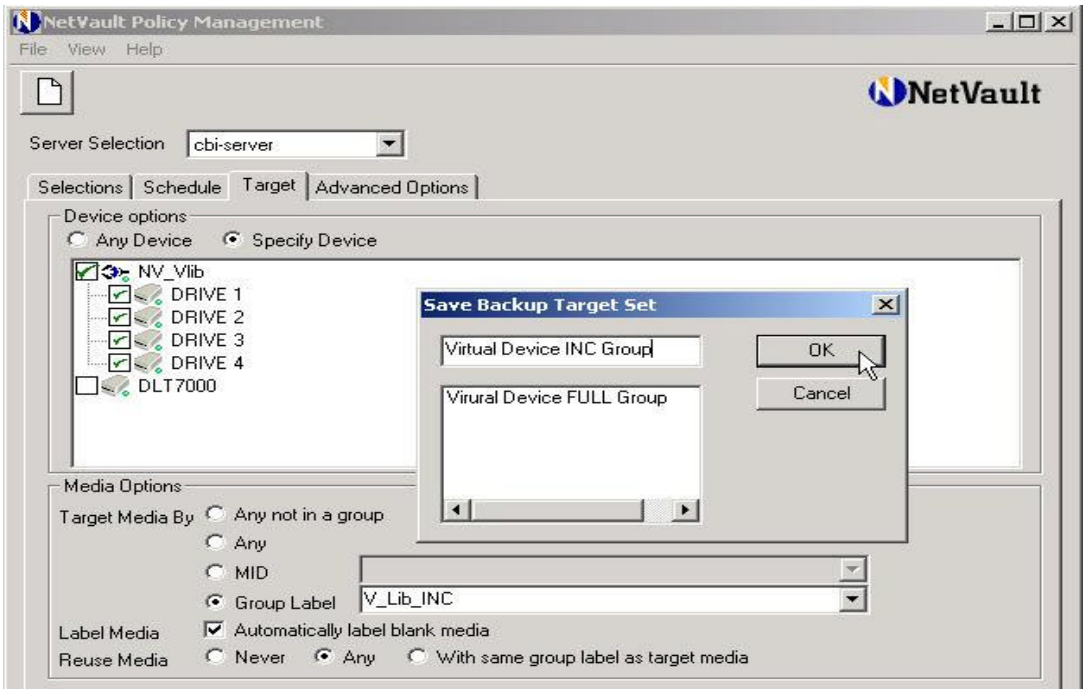

Step 6
Press on the Clear button again; select the Target Tab for our next template.

Step 7
The third Target template we will create will send data to the physical drive(s). This will be used in the duplication phase.
1) De-select all virtual drives leaving only the physical drives you want to use ticked.
2) As with any job, select Reuse Media – *Any* in the Media Options.Click on *Save As* at the bottom of the screen. Save the template with a descriptive name such as *Physical_DUP Device Group*.The template for the physical drives is now ready.

Now that we have created the three Target templates needed for FULL, INCREMENTAL and DUPLICATION, we are ready to create the FULL backup job.

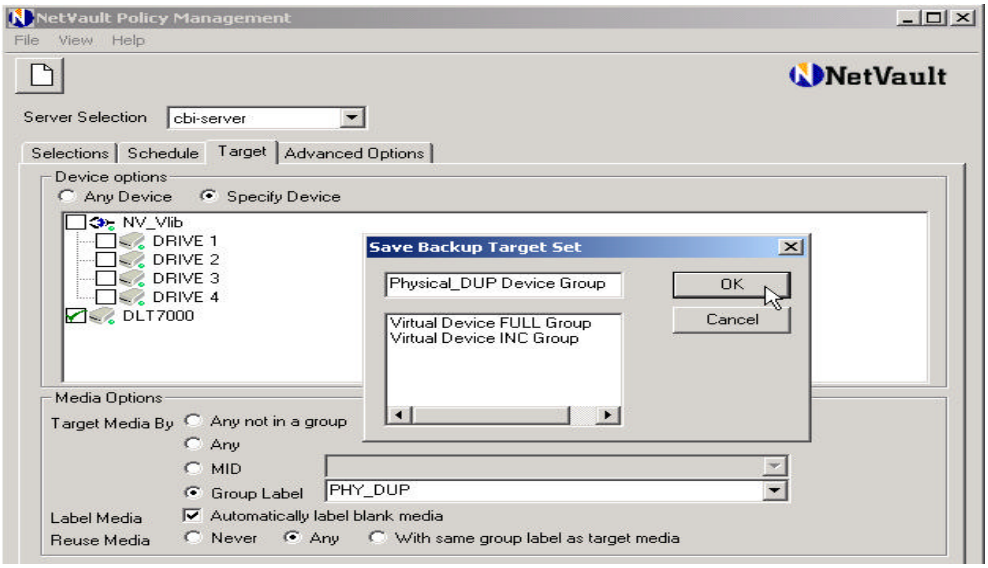

Creating a FULL backup to VDL Staged job

Creating a staged job is very similar to a regular job, the only difference is that you use the Target and Advanced Options to direct and create a two-phase job. Selecting Duplicate under Advanced Options makes it a two-phase job by performing the additional duplicate step.

Step 1

From the Main Screen, select Backup; select the client to be backed up.

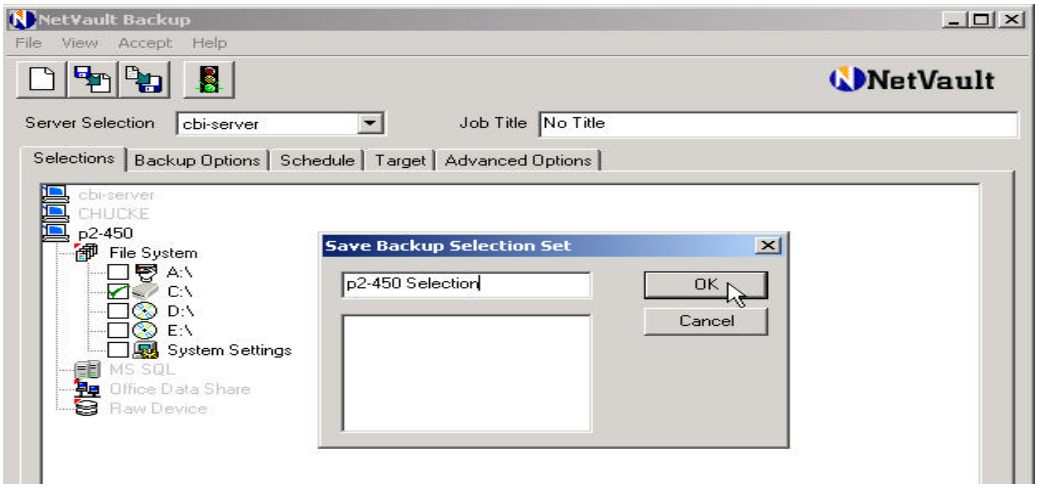

Step 2
Make your Selections, and click *Save as* at the bottom of the screen and save with a descriptive name such as *p2-450 Selection*.

Step 3
Options tab, check *Backup Type*, it is set to Full by default.

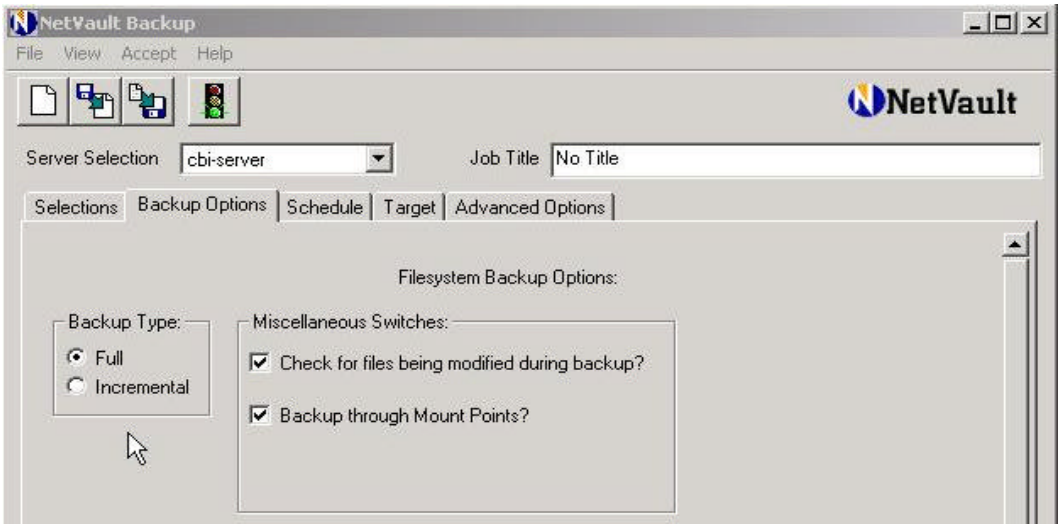

Step 4
Schedule Tab, Schedule as you would a normal job.

Step 5
From the *Target* Tab, Click on the *Load* button on the bottom of the screen and select *Virtual Device FULL Group*.
We want the job to backup to the Virtual library first.

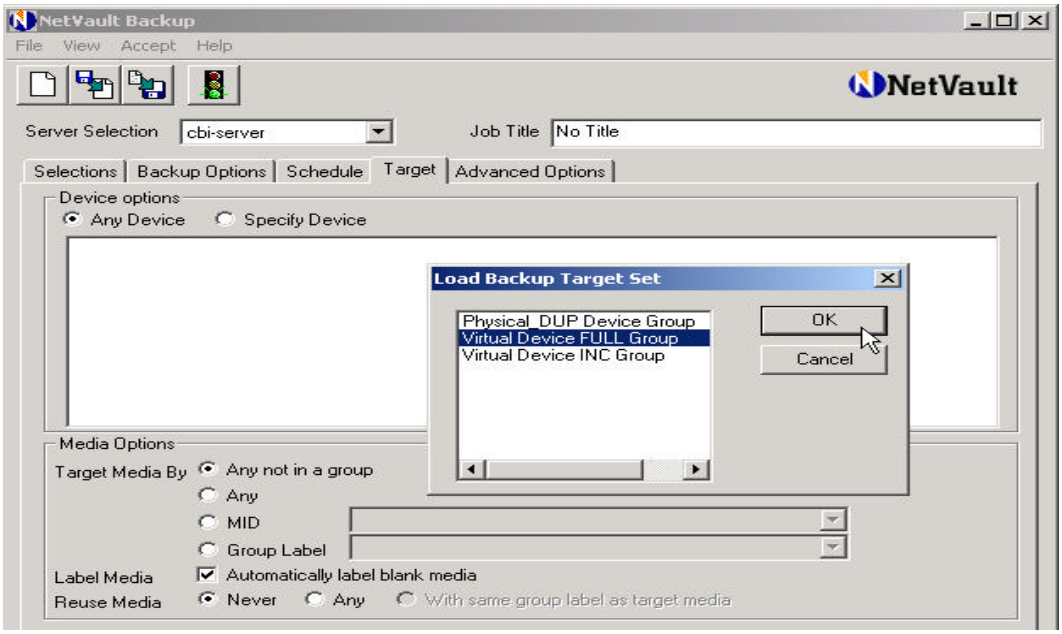

Step 6

From the Advanced Options Tab.

1) Select Duplicate.
   Instructs NetVault to run the second phase, which is to duplicate the backup data on the VDL to tape.
2) Select Migrate (Discard original backup).
   Once the duplication phase 2 completes, NetVault will mark the backup saveset on the Virtual Library as reusable thus freeing up the media for the next backup job.
   Note: The Migrate function is the most common setting, but for these examples we will not be using the Migrate function, instead we are going to use the Discard After function.
   From the Using Target Set list box, select our Physical_DUP Device Group template.
   To direct data copy onto the physical drives.

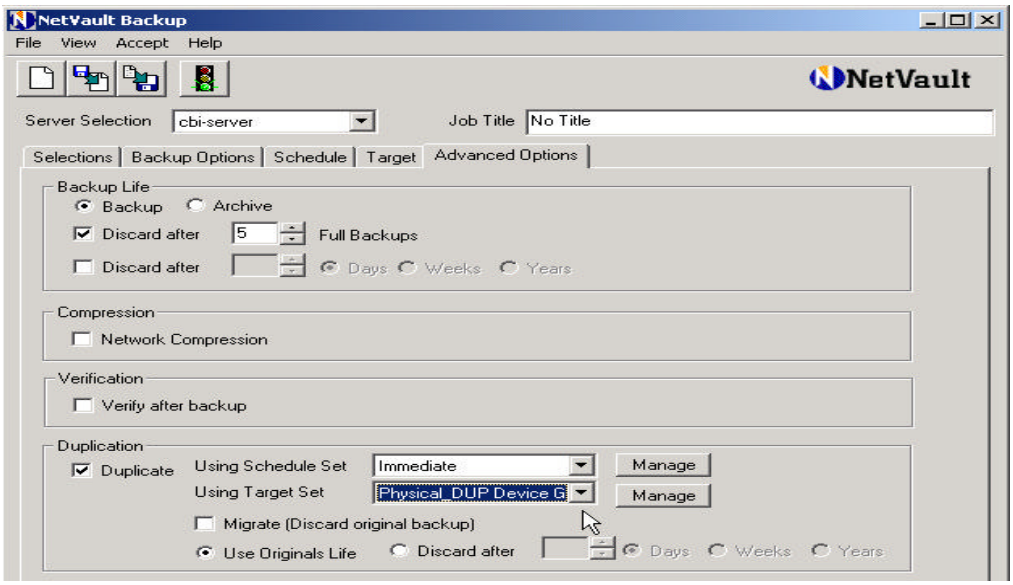

Step 7
Submit the job, with a descriptive name, in the *Job Title* box such as *p2-450 FULL Staged*.

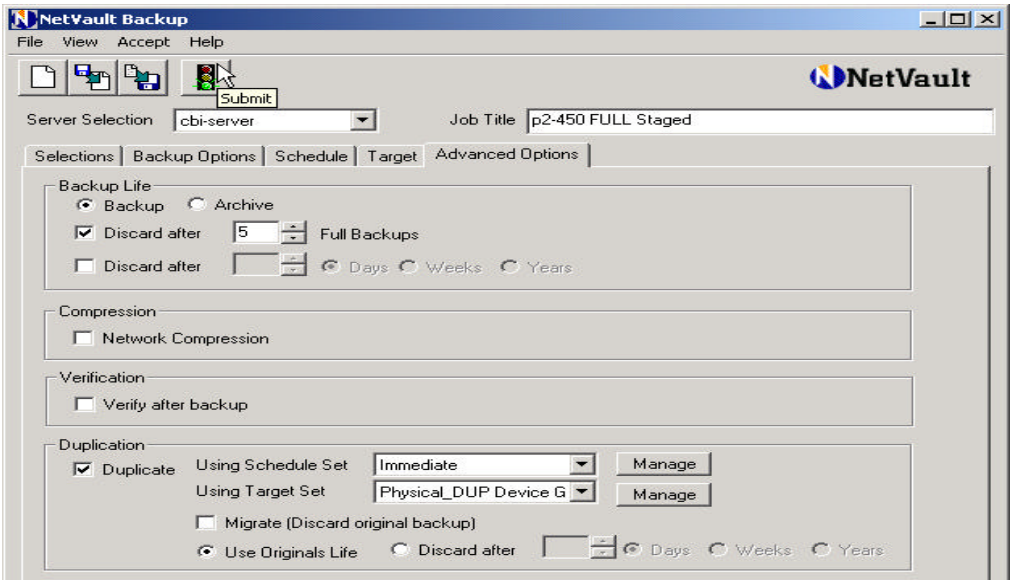

Step 8
After the job completes, you can verify that the data was written how and where you wanted it, by checking both the *Device Management* and the *Media Management* windows. You will see both the V_Lib_FULL and PHY_DUP media groups in each window.

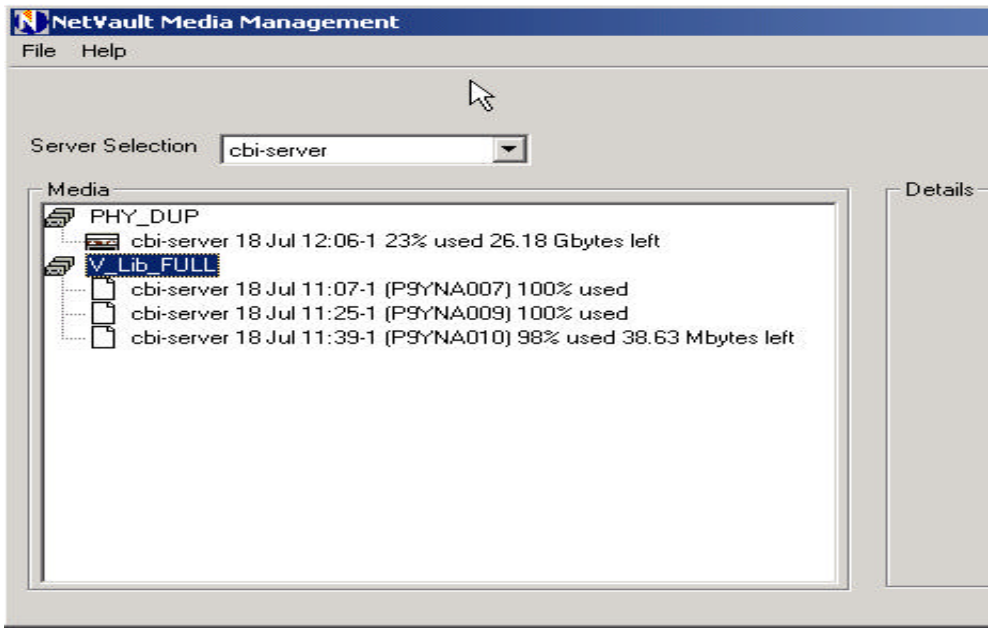

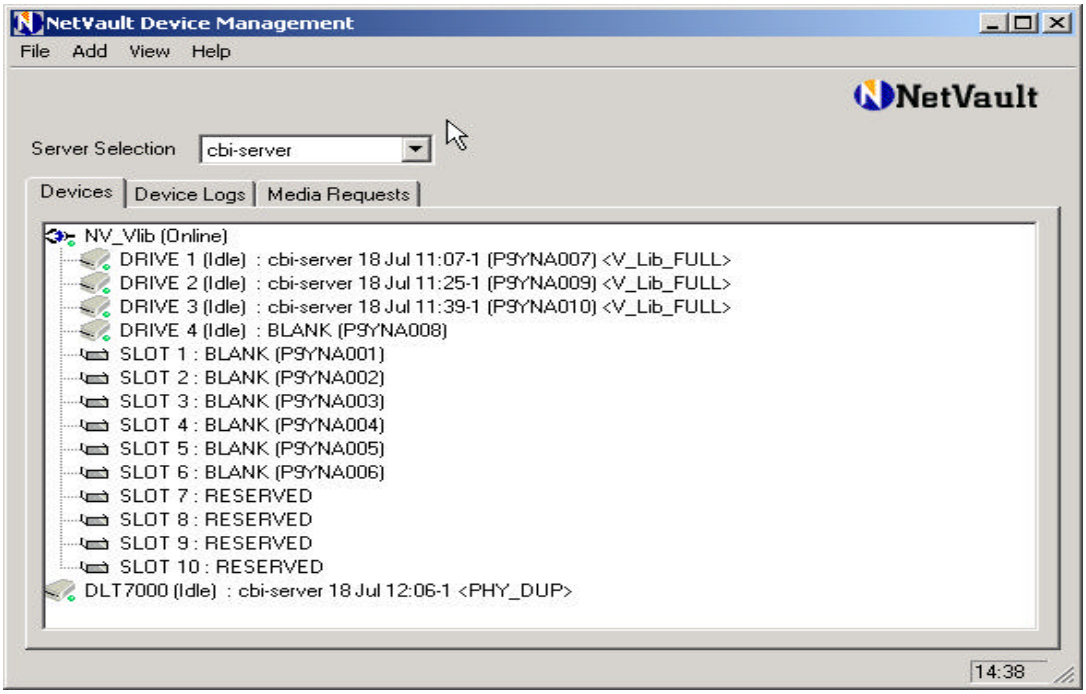

Creating a Incremental VDL Staged Job

Creating an Incremental backup is almost the same as the Full backup job.

Step 1

From the *Selection Tab*, Click on the *Load* button on the bottom of the screen and select *p2-450 Selection*.

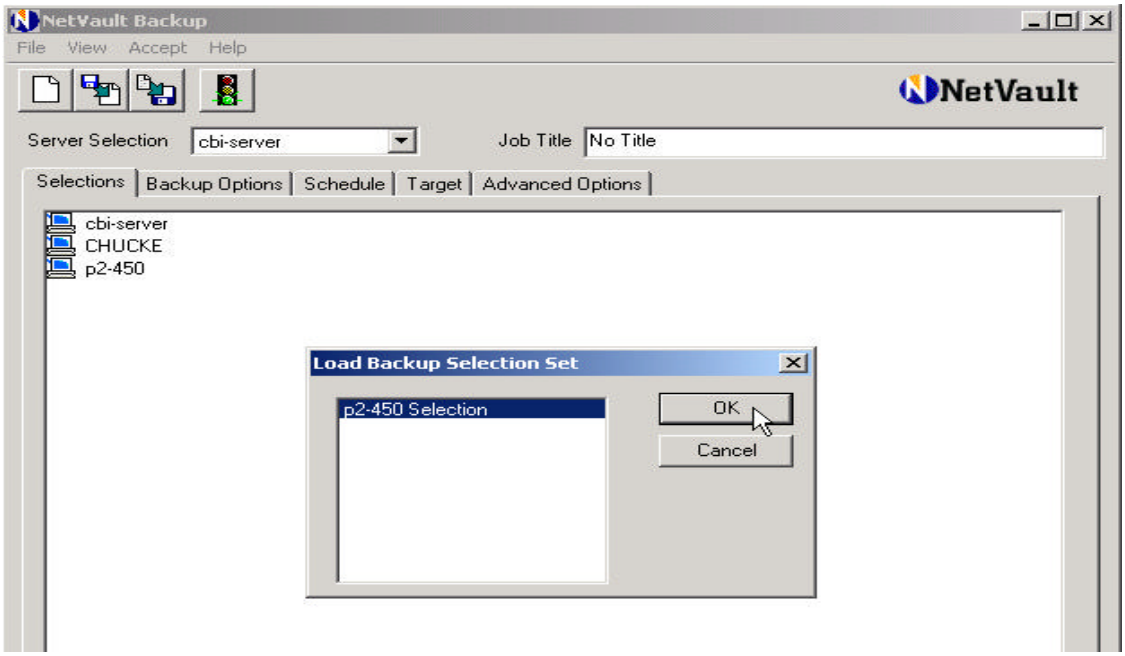

Step 2
From the *Option Tab*, Click on the *Incremental* button on the left side of the screen.

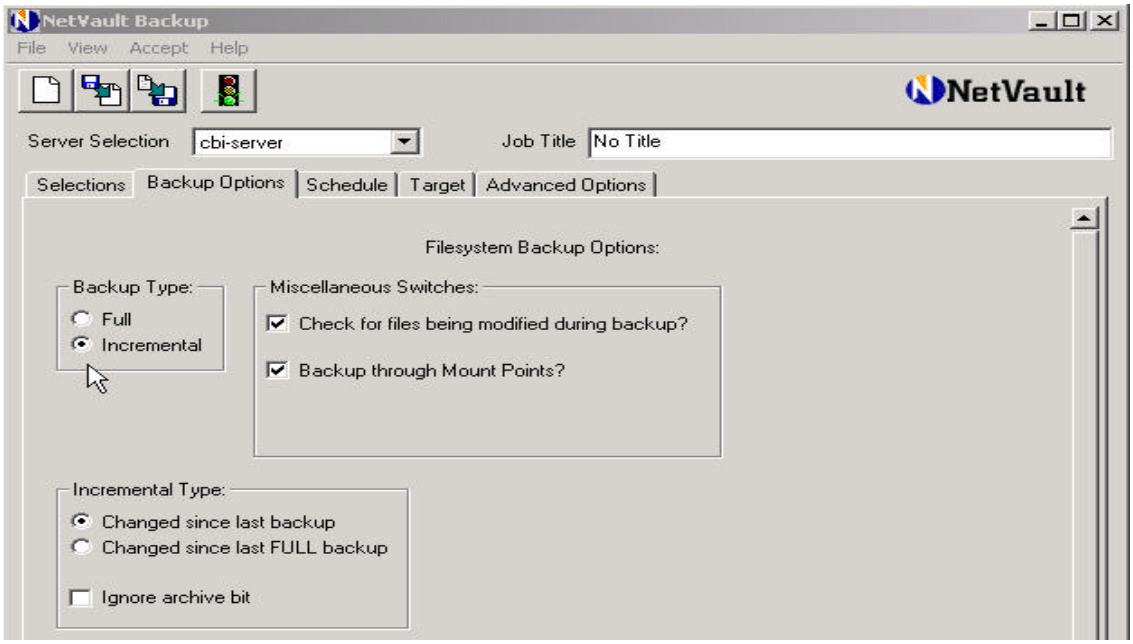

Step 3
From the *Target* Tab, Click on the *Load* button on the bottom of the screen and select *VirtualDevice INC Group*. Again we want the job to backup to the Virtual library first. But this time it is going to a different Target group in VDL.

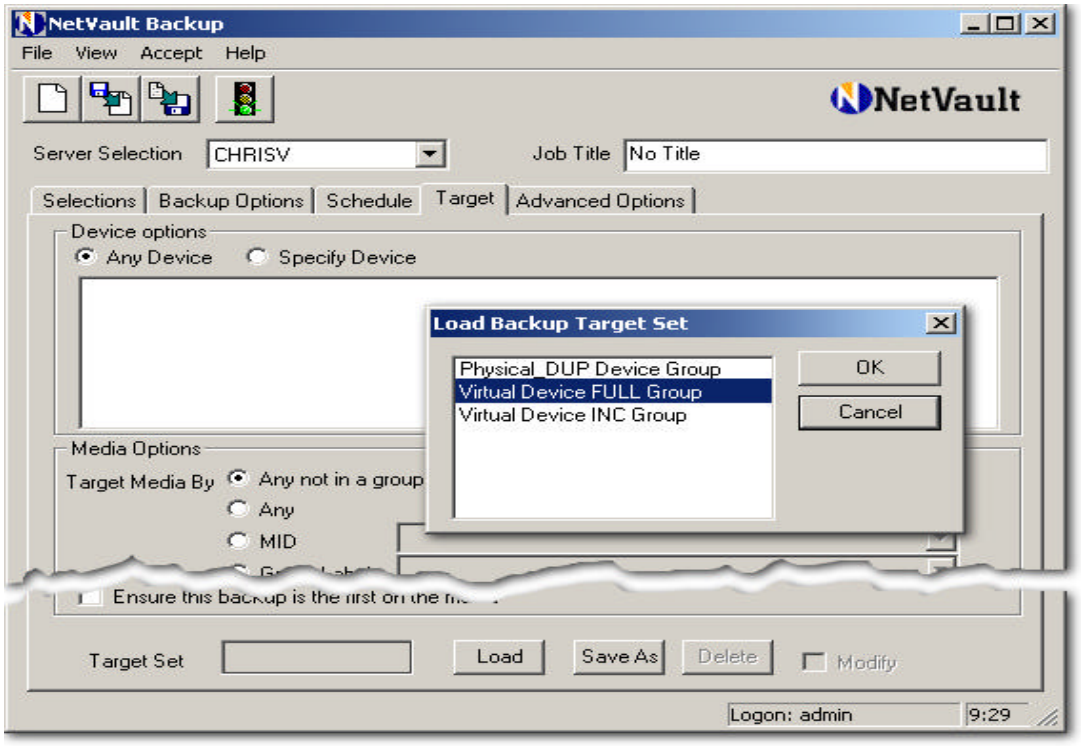

Step 4

From the Advanced Options Tab, under Duplication:

1) Select *Duplicate.*
   Instructs NetVault to run a second phase data copy job
2) From the *Using Selection Set* list box, select our *Physical_DUP Device Group* template.
   Perform the data copy again onto the physical drives.
3) Submit the job, with a descriptive name, in the ***Job** Title* box such as *p2-450 INC Staged.*

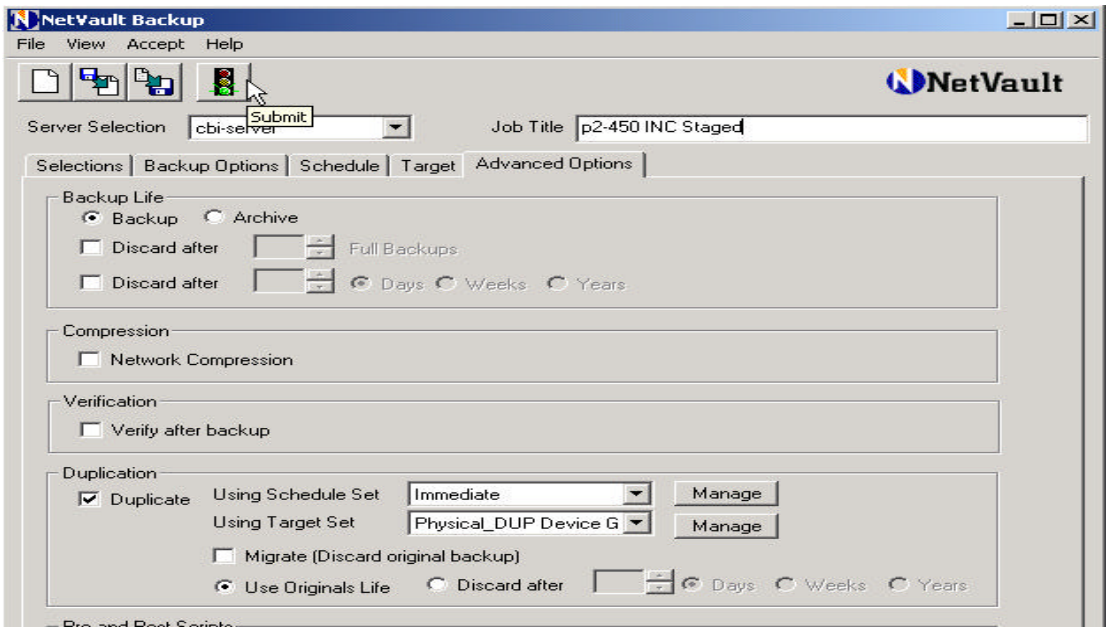

Consolidated Backups

The setup for a consolidated backup is only slightly different than a standard backup. The difference is in the *Selection Tab*. Instead of selecting the Client to be backed up, you need to select the NetVault server; in this case it is *cbi-server*.

Because the consolidated backup is a manipulation of the NetVault database index to create a new Full based on the Full backup and the following Incrementals, you will find the Consolidate File System backup under the NetVault Server.

But first, for the Consolidated backup we need to create one more Target template to complete this backup strategy.

Step 1

1) De-select all physical drives leaving only the virtual drives ticked.
2) To create separate pool of media, enter *Group Label*, V_Lib_CON_FULL
3) As with any job, first select Reuse Media – *Any* in the Media Options.
4) Click on *Save As* at the bottom of the screen. Save the template with a descriptive name such as *Virtual Device CON_FULL Group*.

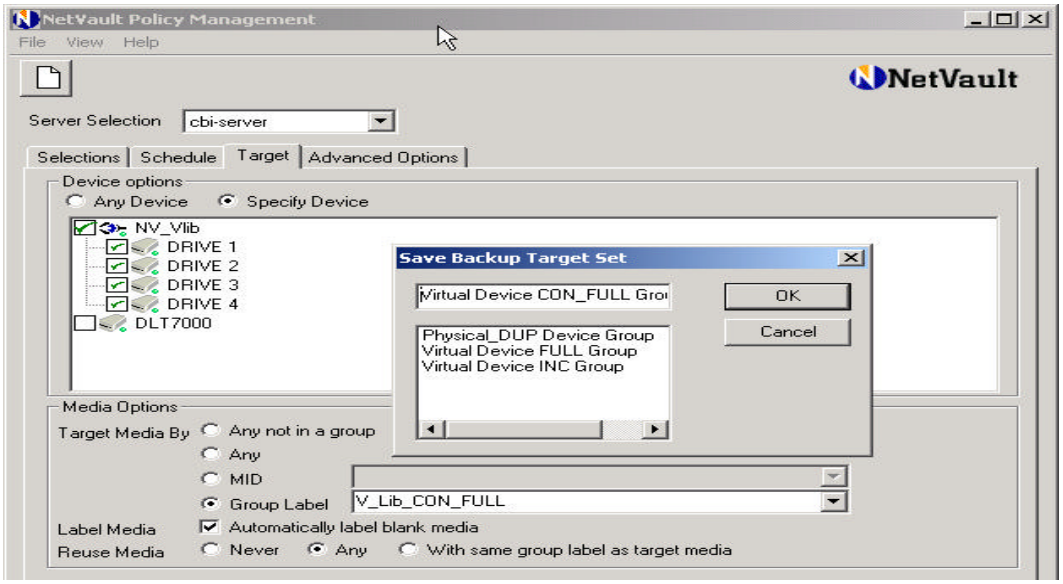

Step 2

From the Main Screen, select Backup, and select the NetVault server.
Select Consolidate Files System Backup

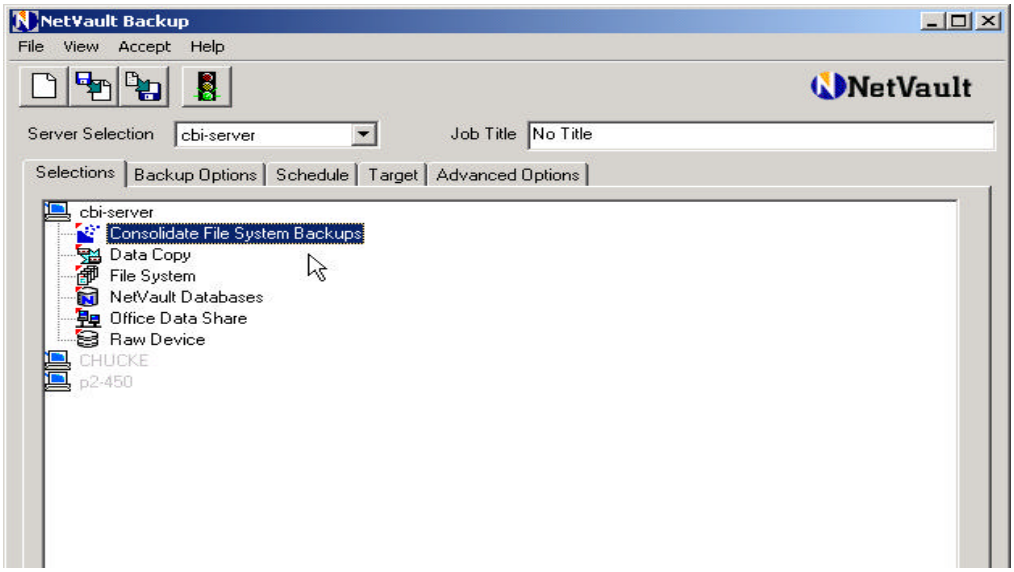

Step 3
Select Filesystem Backup Sets.
This is done in order to use the saved selection set for p2-450. Select it.

Note: There is nothing to set in Backup Options for a Consolidated backup.

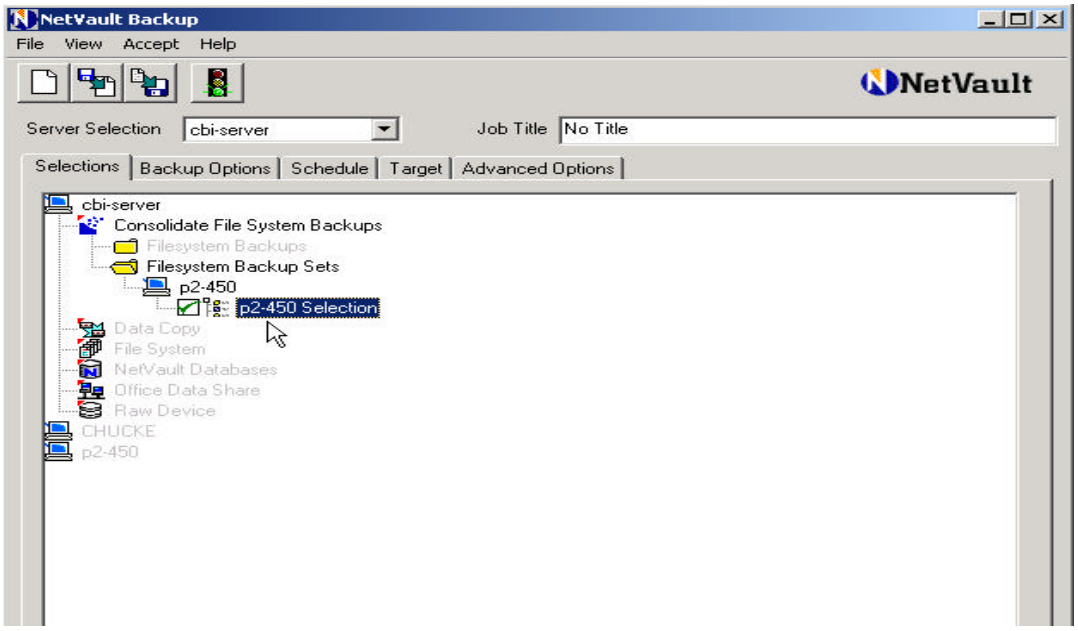

Step 4
From the *Target* Tab, Click on the *Load* button on the bottom of the screen and select *Virtual Device CON_FULL Group*.
We want the job to Consolidate to the Virtual library first. And it is going to a different group.

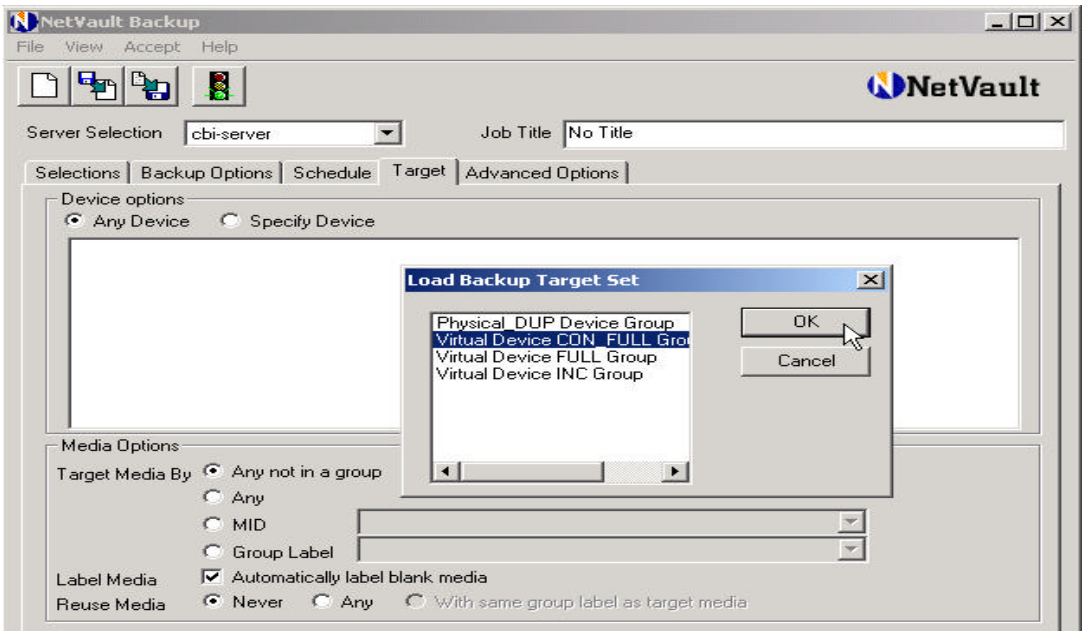

Step 5
From the Advanced Options Tab, under Duplication:

1) Select *Duplicate.*
   Instructs NetVault to run a second phase data copy job
2) From the *Using Selection Set* list box, select our *Physical_DUP Device Group* template.
   Perform the data copy again onto the physical drives.
3) Submit the job, with a descriptive name, in the ***Job** Title* box such as *p2-450 Consolidated Staged.*

Note: in all of the examples, where we have saved names for templates, media pools, and jobs you should create a naming convention, which applies, to your environment.

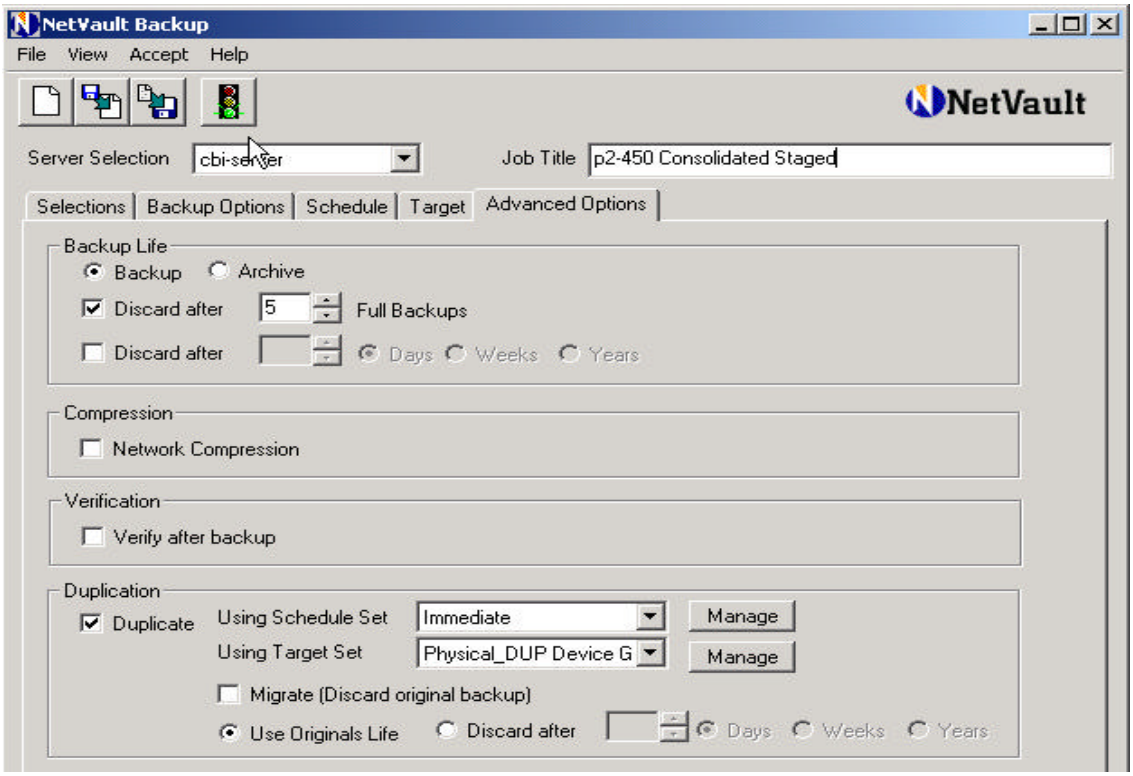

Conclusion

NetVault's Virtual Disk Library is a powerful and flexible tool to help decrease your backup window and at the same time increase your ROI on Tape Libraries by streaming the tape drives at highest rates possible.
The example above is only one strategy using your VDL. There are numerous ways to utilize it, in combination with Virtual and Physical devices.
Please contact BakBone Software Customer Support at support@bakbone.com for further assistance if required.